

A Fair Admission Control Mechanism for Efficient Utilization of Resources in On-chip Nanophotonic Crossbars

Seyed H. Mirsadeghi^a, Ahmad Khonsari^{b,c}, M. Sadegh Talebi^d, and Behnam Khodabandelo^c

^a*Department of Electrical and Computer Engineering, Queen's University, Kingston, ON, Canada*

^b*Department of Electrical and Computer Engineering, University of Tehran, Tehran, Iran*

^c*School of Computer Science, Institute for Research in Fundamental Sciences (IPM), Tehran, Iran*

^d*School of Electrical Engineering, The Royal Institute of Technology (KTH), Stockholm, Sweden*

Abstract

Advances in CMOS-compatible photonic elements have made it plausible to exploit nanophotonic communications to overcome the limitations of traditional NoCs. Amongst various proposed nanophotonic architectures, optical crossbars have been shown to provide high performance in terms of bandwidth and latency. In general, optical crossbars provide a vast volume of network resources that are shared among all the cores within the chip. In this paper, we present a fair and efficient admission control mechanism for shared wavelengths and buffer space in optical crossbars. We model buffer management and wavelength assignment as a utility-based convex optimization problem, whose solution determines the admission control policy. Thanks to efficient convex optimization techniques, we obtain the globally optimal solution of the admission control optimization problem by using simple and yet efficient iterative algorithms. We cast our solution procedure as an iterative algorithm to be implemented a central admission controller. Our experimental results corroborate the gain that can be obtained by using such an admission controller to manage the shared resources of the system. Furthermore, they confirm that the proposed admission control algorithm works well for various traffic patterns and parameters, and evinces a tractable scalability with increase in the number of cores of the crossbar.

Key words: Convex optimization, Fairness, Nanophotonic crossbar, Optical chip multiprocessor, Wavelength assignment.

Email addresses: s.mirsadeghi@queensu.ca (Seyed H. Mirsadeghi), ak@ipm.ir (Ahmad Khonsari), mstms@kth.se (M. Sadegh Talebi), bk@ipm.ir (and Behnam Khodabandelo)

1. Introduction

According to Moore's law [1], the number of transistors on a single chip doubles every two years and this trend has made it plausible to integrate transistors on a single chip in the scale of billions. Current trend for exploiting such a large number of transistors is to have multiple computation cores on the chip rather than a single powerful processor, taking advantage of parallel computing [2]-[3]. This has led to Chip MultiProcessor (CMP) and System-on-Chip (SoC) design paradigms with an ever increasing number of cores. In such paradigms, the communication among cores is a key design factor, which is considered to be the performance bottleneck for many-core CMPs. In fact, the communication bottleneck has made a shift in chip design from a computation-centric approach to a communication-centric one [4, 5].

Network-on-Chip (NoC) was proposed as a promising design alternative to satisfy on-chip communication demands [6, 7]. A NoC design uses a packet-switched network comprising a set of routing elements connected to each other by a set of links. Despite their success in improving communication performance, traditional NoCs designed based on copper wires have been facing various challenges towards a scalable solution [8]. In the billion transistors era, wires do not scale as they occupy too much of area, and are difficult to route within the chip. In addition to electromagnetic interference problems such as the wire crosstalk, wires contribute to a large portion of the chip power consumption. Moreover, high latency in transferring messages between far nodes¹ and the lack of sufficient bandwidth add even more limitations to traditional NoCs [9].

In accordance to the abovementioned limitations, the nanophotonic communication paradigm has gained a lot of interest as an alternative to traditional copper-based NoCs. In nanophotonic communication, data is transferred by means of light over an on-chip optical medium. Although the idea was initially proposed four decades ago [10], practical limitations in optical device integration at that time put its usage off until recent years. Thanks to advances in CMOS-compatible photonic elements, optical interconnects are now considered as a practical choice for CMPs communication infrastructure [11, 12, 13].

Utilizing optical interconnect can remedy many of on-chip communication restrictions faced by electrical interconnects such as crosstalk, voltage isolation, wave reflection, and etc. [14]. Moreover, the high propagation speed of light in optical waveguides along with wavelength division multiplexing (WDM) techniques result in very high bandwidth and low latency communications in optical interconnects. Additionally, optical interconnects have the advantage of bit-rate transparency and low loss of optical waveguides that can

¹Throughout the rest of this paper, we use the terms 'core' and 'node' interchangeably.

highly reduce power consumption compared to their copper-wire counterparts [15, 16].

Various architectures have been proposed to exploit nanophotonic technology for on-chip communication networks. Among them, taking advantage of WDM techniques, optical crossbars have been shown to provide high performance in terms of bandwidth and latency. In addition, compared to other architectures, optical crossbars generally provide a vast volume of network resources shared among cores. On the other hand, NoC architectures are expected to provide different service levels to support a variety of application requirements in SoC and CMP designs. Generally speaking, in a NoC where multiple applications compete to access shared resources, a mechanism is needed to realize fair allocation of resources while guaranteeing some predefined quality of service (QoS) requirements. To this end, several researchers have studied QoS provisioning or service differentiation schemes in traditional NoCs [17, 18, 19].

The aim of this paper is to present an admission control mechanism for fair and efficient allocation of shared resources in an optical crossbar. Specifically, we consider two major resources in an optical crossbar: a) the set of wavelengths, and b) the end nodes' buffer space. Accordingly, we model wavelength assignment and buffer management in an optical crossbar as a utility-based admission control problem, using the Network Utility Maximization (NUM) framework for resource allocation in data networks [20]. First, we model wavelength assignment and buffer management in the crossbar as a rate control scheme. Next, we formulate the admission control policy as a rate allocation optimization problem. We consider concave utilities corresponding to applications with elastic traffic, and therefore model the admission control as a convex optimization problem. Thanks to the well-established theory of convex optimization, we obtain the globally optimal solution of the problem by using simple and yet efficient iterative algorithms. We then cast the solution as an iterative admission control algorithm to be implemented by a central controller that will be in charge of running the algorithm. Based on the results of the admission control algorithm, the central controller determines which source node may send data over each wavelength. Our simulation results confirm that the proposed admission control algorithm works well for various traffic patterns, and evinces a tractable scalability with increase in the size of the crossbar. Moreover, the results corroborate that the use of the proposed controller to manage shared system resources could delicately balance performance and fairness.

The rest of the paper is organized as follows. In Section 2, we review the related work. Section 3 describes the underlying architecture, followed by the system model description in Section 4. Section 5 is devoted to cast the admission control procedure as the solution to a convex optimization problem. Section 6 investigates the optimal solution, followed by the corresponding algorithms in Section 7. Simulation results

are reported in Section 9. Finally, Section 10 concludes the paper and outlines some future directions.

2. Related Work

Various approaches have been proposed to exploit nanophotonics for on-chip communication networks. Development of efficient optical data buffers is still a challenging problem, which makes it quite hard to construct a fully optical packet-switched network. As a result, Shacham et al. [16] propose a hybrid approach, in which large data packets are transferred using an optical circuit-switched network whereas the optical network is controlled by an electronic packet-switched network. Adi et al. [21] consider the architecture proposed in [16] and propose a predictive switching and reservation technique to reduce its path setup latency. In some other hybrid architectures such as [22], local communications are carried out by an electrical network, whereas long distance communications utilize optical links. Ahmed et al. [23] exploit non-blocking photonic switches as well as light-weight electronic routers to decrease the latency and power overheads of hybrid architectures. García-Guirado et al. [24] propose a set of policies to manage hybrid networks consisting of ring-based photonic and electrical mesh sub-networks. The proposed policies use different criteria, such as message size, distance, and photonic ring availability, to decide which sub-network to be used for each message.

ATAC [25] is another hybrid architecture in which a baseline electrical 2D mesh is used for close-range point-to-point communications, whereas a ring-like optical network is used for long-distance and collective communications. The optical interconnect functions in a similar way as a broadcast bus, and contention is resolved by assigning unique wavelengths to senders. SUOR, proposed by Wu et al. [26], uses a circuit-switched ring-based optical NoC with a control sub-system responsible for arbitration and flow control. The control sub-system sets up the path from source to destination based on the requests received from nodes. SOUR also takes advantage of channel segmentation by dividing one waveguide into multiple non-overlapping sections that can support multiple transactions simultaneously.

Briere et al. [27] use 4-port optical switches to build a photonic routing structure called λ -router, which provides contention-free communication among cores through wavelength routing techniques. Each pair of cores communicate through fixed and predefined wavelengths routed passively by the λ -router. *CoNoC* is the architecture proposed by Koohi et al. [28], where all-optical switches are used to implement contention-free wavelength-based passive routing of optical streams. Contention-free communication is carried out by assigning each node with a unique wavelength for data reception. This way, contention is confined to the end-points and is resolved by an electrical arbitration scheme. A scalable wavelength-routed optical NoC

based on the Spidergon topology is presented in [29]. It uses per-receiver wavelengths in the data network to prevent network contention, and adopts per-sender wavelengths in the control network to avoid end-point contention. Werner et al. [30] also propose an all-optical design exploiting a mesh-like topology.

The architecture proposed by Kirman et al. [31] is an instance of an optical crossbar where several Single Write, Multiple Read (SWMR) busses provide full connection among nodes. Each SWMR bus is dedicated exclusively to one node for sending data to others, whereas all other nodes can read data from all busses. A comparison study of worst-case optical losses for various crossbar implementations is presented in [32]. Pan et al. [33] introduce *Firefly*, which is a hybrid crossbar with nodes partitioned into clusters. While intra-cluster communication is done using smaller electrical crossbars, inter-cluster communication is realized by a SWMR optical crossbar. Unlike the architecture proposed in [31], optical packets are not broadcast to all nodes, but rather, the intended receiver is selected by auxiliary *reservation channels* prior to each communication.

Corona is an all-optical crossbar topology, proposed by Vantrease et al. [34], which implements a crossbar consisting of Multiple Write, Single Read (MWSR) shared busses. Each bus is dedicated exclusively to one node for receiving data, whereas all nodes can write data on all busses. Corona takes advantage of an optical token-based arbitration mechanism to resolve contention among nodes for sending data on the same bus. To address the fairness issues, a fair token slot mechanism is presented in [35]. Fu et al. [36] employ MWSR token-ring busses of Corona to implement rows and columns of an optical 2D torus, and modified the token-based arbitration scheme so as to support virtual channel flow control.

In [37], frame-based arbitration is used to provide QoS, in terms of differentiated bandwidth allocation, for an architecture similar to Corona. The so-called *FlexiShare* architecture presented in [38] can be viewed as a combination of Firefly and Corona. FlexiShare makes use of Multiple Write, Multiple Read (MWMR) shared busses, where each node may write data on, or read data from, any of the shared busses. It has the advantage of lower power consumption and better channel utilization compared to other optical crossbars. Li et al. [39] propose *LumiNoC* in which the network is broken into several smaller subnets so as to avoid long waveguides. All nodes in the same subnet are connected to the same waveguide. However, communication between two subnets requires a hop through an intermediate electrical router.

Closest to our work is the work done by Pan et al., which present *FeatherWeight* [40]. FeatherWeight provides an optical arbitration scheme with QoS support in nanophotonic MWSR crossbars. Similarly to [35], they use token streams to grant access to source nodes for sending data to each home (destination) node. In addition, each source node is assigned a “quota” that designates the maximum number of tokens a

node can grab in each time slot. To improve resource utilization and enforce QoS, the quotas are dynamically changed with respect to the request patterns. A QoS controller residing at each home node is used to update the quotas. At each time slot, every node gives feedback to the controller indicating the amount of tokens it has consumed in the previous time slot. Having gathered the feedback, the controller first updates the quotas and then propagates the updated values to the nodes. Access to data channels is granted to nodes with respect to the new quotas in the following time slot. The proposed QoS mechanism provides both fairness and differentiated service among the nodes.

Similarly to FeatherWeight, we propose a time-slotted mechanism to enable fairness and differentiated service among the nodes of a nanophotonic crossbar. However, FeatherWeight is based on an MWSR architecture, whereas we target MWMR photonic crossbars. In FeatherWeight, each MWSR channel is arbitrated separately, and fairness is enforced on each MWSR channel independently of other MWSR channels. We present an admission control scheme that takes into account all available channels in the optical crossbar as a whole, and provides fairness and QoS globally among all nodes and all data channels. In contrast to FeatherWeight that only provides max-min fairness, our channel allocation scheme could cover a larger class of fairness metrics. Finally, we use a centralized controller to carry out resource allocation among the nodes. We use a different QoS algorithm that models the problem as an optimization problem whose solution is the optimal resource allocation in each time slot.

3. Nanophotonic Architecture

Three basic blocks are needed to realize a nanophotonic interconnect: *silicon waveguides*, *microring resonators*, and *laser sources*. Waveguides are the medium over which the optical signals are transferred at the speed of light. Using DWDM² techniques, multiple wavelengths can travel in a single waveguide simultaneously without interfering with each other. Microring resonators are the dominant state-of-the-art elements used in optical NoCs for modulation and/or detection of particular wavelengths. A light generation source is also required to provide the beam of light over which data is modulated. Due to difficulties in integrating a silicon-based laser onto a chip, off-chip laser sources are preferred. The light generated off-chip is coupled onto the chip by means of optical fibers. Figure 1 shows a basic configuration of nanophotonic elements for optical on-chip communications.

²Dynamic Wavelength Division Multiplexing




Figure 1: Basic configuration of nanophotonic elements for optical data transmission

The architecture we consider in this work is similar to the one proposed in [38] referred to as Flexishare. As mentioned in Section 2, it uses Multiple Write, Multiple Read (MWMR) shared busses to implement a crossbar among the nodes, and each node may write or read data to/from any of the shared busses. In such an architecture, the whole network can be considered as a pool of communication channels shared among all nodes that can be dynamically configured to support data transmission between multiple source-destination pairs. From nanophotonic point of view, channels are realized through several waveguides deployed on the chip. We consider the same waveguide layout presented in [41], where each waveguide starts its path from the upper left corner of the chip and visits all nodes along a serpentine path that terminates at the upper right corner. Figure 2(a) depicts a schematic diagram of such a layout for a 4×8 NoC.

Terminating waveguides at the upper right corner leads to a *single-round* implementation of data channels as each waveguide passes all nodes exactly once. To provide full connectivity, each node needs to modulate light in opposite directions on the waveguides. Therefore, the channels should be divided into two sets, where in one set light is injected from the upper left corner of Figure 2(a) propagating toward the upper right corner, and vice versa for the other set. In Figure 2(a), the two sets have been shown in blue and red, respectively. Depending on the relative position of the source and destination nodes (which dictates the required direction of light transmission), channels from the first or the second set will be used.

In Figure 2(a), all nodes have been indexed from 1 to 32. An arbitrary source node with index i can use the set of channels shown in blue for sending data to destination node j only if $i < j$. Similarly, red channels can be used only when $i > j$. Consequently, though a single-round layout of waveguides benefits from shorter length and lower power dissipation, it imposes some limitation in terms of channel sharing

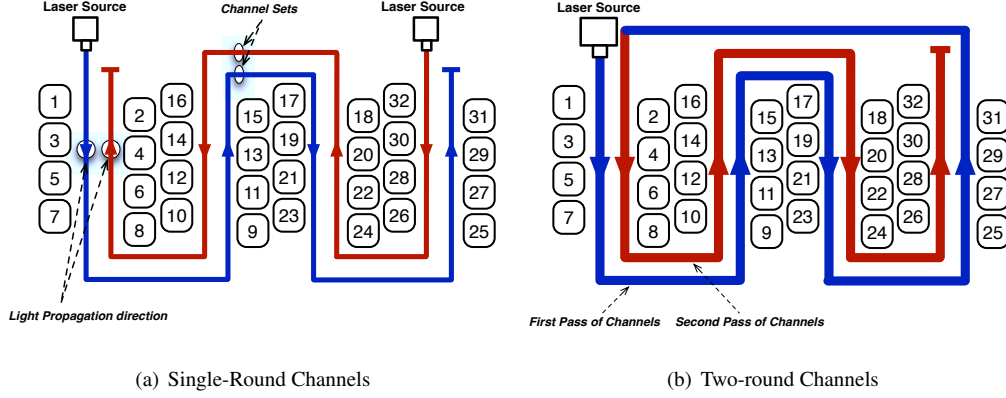


Figure 2: Layout of waveguides

because the channels in one set cannot be utilized for data transmission in the opposite direction. To clarify this, consider the scenario in which nodes $1, 2, \dots, 31$ wish to send data to nodes $2, 3, \dots, 32$, respectively. Since each source node has a lower index than its corresponding destination node, all the senders can only utilize the set of blue channels for data transmission. This results in a high contention over the channels in one set, as well as a low utilization of the channels in the other set.

In order to mitigate this problem, which limits our ultimate goal of total sharing of channels, we use another implementation of optical data channels called *two-round* channels [38]. The waveguide layout is the same as the one shown in Figure 2(a), but each waveguide continues its path back to the originating point at the upper left corner rather than being terminated at the upper right corner. Figure 2(b) portrays an instance of such a layout. In this way, each waveguide will pass each node twice along its path, and the light will only be injected in one direction. For all nodes, the first pass of each waveguide (shown in blue in Figure 2(b)) may only be used for light modulation, whereas detection takes place only on the second pass (shown in red). Thus, a node can utilize any of the available channels across all waveguides for data transmission to any other node irrespective of the relative position of the sender and receiver. With the same number of waveguides, a two-pass layout will potentially provide each node with more channels for data transmission compared to a single-pass layout. This is shown by using thicker lines in Figure 2(b) to illustrate the waveguides.

With the above-mentioned design, each node can potentially make use of all the channels for data reception and/or transmission. This brings a higher flexibility in terms of channel sharing among the nodes, which in turn provides better utilization of channels. A major challenge in such an architecture will be the efficient

management and allocation of channels. We will deeply explore this issue in the following sections.

4. System Model

We consider the architecture described in Section 3 with N nodes indexed by $n \in \{1, \dots, N\}$, and W waveguides deployed in a two-round fashion. We assume 64 wavelengths are multiplexed on each waveguide providing $C = 64 \times W$ channels in the system. Letting R denote the data rate of an individual wavelength, the transmission rates will be multiples of R . In order to transmit data, each sending node should be allocated at least one of the C wavelengths. Allocating more wavelengths to a node provides it with a higher rate for data transmission. In addition, each wavelength can be allocated to at most one node at any moment.

We partition time into multiple slots of length δ . Let \mathcal{I}^t denote the set of all nodes that have data for transmission at time slot t . We define $\mathcal{I}_k^t \subset \mathcal{I}^t$ as the set of nodes wishing to send data to destination node k . Equivalently, node k is the destination node for the set \mathcal{I}_k^t . Moreover, we assume that each node will have some specific buffer space for receiving data from other nodes. For any node k , the length of this buffer depends on both the cumulative rate at which other nodes send data to node k , and the drain rate of node k . Drain rate of a node denotes the rate at which a node can process received data. We assume that at time slot t , node k can process incoming packets with rate r_k^t , has M_k^t free buffer space, and receives data from node $n \in \mathcal{I}_k^t$ at rate x_{nk}^t . It is worth mentioning that x_{nk}^t is proportional to the number of wavelengths (channels) assigned to each node for data modulation.

We represent by $\mathbf{x}_n^t = (x_{nk}^t, k = 1, \dots, N)$ the rate allocation vector or simply rate vector for (sender) node n at time t . We also denote the overall rate allocation at time slot t by the vector $\mathbf{X}^t = [\mathbf{x}_1^t, \dots, \mathbf{x}_N^t]$ made by concatenating rate vectors of all (sender) nodes. The rate allocation policy at time slot t is feasible if and only if the corresponding rate allocation vector $\mathbf{X}^t \in \mathbb{R}^{N^2}$ satisfies:

$$x_{nk}^t \geq 0; \quad n, k = 1, \dots, N, \quad (1)$$

$$\sum_{n \in \mathcal{I}_k^t} x_{nk}^t \leq r_k^t + \frac{M_k^t}{\delta}; \quad k = 1, \dots, N, \quad (2)$$

$$\sum_{n \in \mathcal{I}^t} \sum_{k=1, k \neq n}^N x_{nk}^t \leq CR. \quad (3)$$

The first condition requires that all allocated rates be nonnegative. The term $\frac{M_k^t}{\delta}$ in the r.h.s. of (2) represents the rate at which a receiving node can buffer the incoming packets, i.e., the receiver buffering rate. Thus, the buffering rate of a receiver is the rate at which free buffer space is available to its transmitting node(s). The

second equation accounts for feasibility of rate allocation in terms of flow control, and provides an upper bound for the cumulative rate at which nodes can receive data without buffer overflows. This upper bound depends on both the drain and buffering rates. Thus, (2) states that for each receiving node, the cumulative rate of incoming data from all other senders should be no greater than the sum of drain and buffering rates³. The third equation implicitly accounts for feasibility in terms of wavelength availability. It states that the sum of the rates assigned to all nodes for data transmission cannot be greater than the total rate provided by the pool of wavelengths. As long as this condition holds, there will be enough channels available to each node with the assigned rate, provided that each channel is allocated to one sender at most. For the sake of convenience in our derivations in later sections, we introduce *sender-receiver traffic matrix* or simply *traffic matrix* $\mathbf{A}^t = [a_{nk}^t]_{N \times N}$ with $a_{nk}^t = 1$ if $n \in \mathcal{I}_k^t$ and $a_{nk}^t = 0$ otherwise. Indeed, \mathbf{A} represents the traffic pattern of the whole system in terms of the communications between all source-destination pairs.

5. Admission Control Optimization Problem

In this section, we use the rate allocation model discussed in Section 4 to formulate the admission control problem as a rate allocation optimization problem that finds the best rate allocation among all feasible rate vectors in view of (1)-(3). We seek to derive a rate allocation policy which leads to a good channel utilization by using as many available channels (wavelengths) as possible, as well as a fair assignment of the resources to nodes. To this end, we cast the admission control policy as the solution to a utility-based optimization problem that strives to find a rate vector with the highest satisfaction among the set of all feasible rate vectors. In order to quantify the satisfaction level of nodes, we use the notion of utility function. More specifically, let $U_{nk}(\cdot)$ denote the utility function assigned to each logical connection (n, k) . This implies that if at time slot t , node n sends data to node k at rate x_{nk}^t , it will attain a satisfaction level quantified by $U_{nk}(x_{nk}^t)$. We assume that the utility function $U_{nk}(\cdot)$ satisfies the following conditions [42]:

C1: The function $U_{nk}(\cdot)$ is continuous, increasing, and twice differentiable over $(0, \infty)$.

C2: The function $U_{nk}(\cdot)$ is strictly concave with bounded curvature.

It is worth mentioning that the above conditions are not restrictive as our focus is on applications that admit elastic traffic demand. In other words, traffic characteristics of such applications can be efficiently

³Indeed, we assume that each node knows its available buffer space and drain rate a priori. In a practical setting, however, each node might need to predict them to be within some specified error.

captured by a utility function satisfying conditions **C1** and **C2**. A well-known class of utility functions that satisfies conditions **C1** and **C2** is the class of α -fair utility functions introduced in [42], where every utility function is defined based on a fairness parameter $\alpha > 0$ as

$$U(x, \alpha) = \begin{cases} w \frac{x^{1-\alpha}}{1-\alpha} & \alpha \neq 1 \\ w \log x & \alpha = 1 \end{cases}, \quad (4)$$

where w is a positive weight. The choice of α determines the tradeoff between the throughput and fairness. For larger values of α , the system sacrifices the throughput to allocate resources in a *fairer* way. As $\alpha \rightarrow \infty$, we approach the max-min fair allocation, which yields the best fairness at expense of the worst throughput.

Taking into consideration all receiving nodes to which node n may have data to send, we define the utility of node n as the sum of the utility functions for its logical connections as

$$U_n(\mathbf{x}'_n) = \sum_{k=1}^N a'_{nk} U_{nk}(x'_{nk}).$$

We then define the utility of the whole crossbar as

$$U(\mathbf{X}') = \sum_{n=1}^N U_n(\mathbf{x}'_n) = \sum_{n=1}^N \sum_{k=1}^N a'_{nk} U_{nk}(x'_{nk}). \quad (5)$$

In the rest of this paper, we focus on the time slot t , and hence omit the superscript t hereafter. Following the Network Utility Maximization (NUM) framework (see, e.g., [20]), our objective is to assign available wavelengths to senders so as to maximize the total utility, defined in (5), over all feasible rate vectors:

$$\max_{\mathbf{x} \geq 0} \sum_n \sum_k a_{nk} U_{nk}(x_{nk}) \quad (6)$$

$$\text{subject to : } \sum_{n=1}^N a_{nk} x_{nk} \leq r_k + \frac{M_k}{\delta}; \quad k = 1, \dots, N, \quad (7)$$

$$\sum_{n=1}^N \sum_{k=1}^N a_{nk} x_{nk} \leq CR. \quad (8)$$

The optimization variable in (6) is the vector \mathbf{X} . We denote the optimal rate vector for the above problem by $\mathbf{X}^* = [\mathbf{x}_1^*, \dots, \mathbf{x}_N^*]$, where $\mathbf{x}_n^* = (x_{nk}^*, k = 1, \dots, N)$, for $n = 1, \dots, N$. The optimal rate vector determines the optimal admission control policy. Due to conditions **C1-C2**, the objective function of problem (6)-(8) is strictly concave as it is a nonnegative sum of strictly concave functions. Moreover, its constraints are affine functions. Thus, problem (6)-(8) is strictly convex [43]. Furthermore, the feasible region, i.e., the polyhedron defined by constraints (7)-(8), is connected and bounded, and hence is a compact set. Thus, at

least one optimal solution exists. Finally, the maximizer is unique due to the strict convexity of the problem [43].

6. Optimal Solution

This section is devoted to solving problem (6)-(8). The standard approach to solve such a constrained convex problem is to solve the (unconstrained) dual of problem . Duality theory guarantees that under mild conditions, which hold for our problem, solving the dual yields the optimal solution to the primal problem. To do to, in what follows we first establish its Lagrangian and derive the dual function. We then formulate the dual problem and solve it using iterative methods.

6.1. Primal Optimality Analysis

We start by writing the Lagrangian for problem (6)-(8) [44]:

$$L(\mathbf{X}, \boldsymbol{\lambda}) = \sum_n \sum_k a_{nk} U_{nk}(x_{nk}) - \lambda_0 \left(\sum_{n=1}^N \sum_{k=1}^N a_{nk} x_{nk} - CR \right) - \sum_{k=1}^N \lambda_k \left(\sum_{n=1}^N a_{nk} x_{nk} - r_k - \frac{M_k}{\delta} \right). \quad (9)$$

Here, $\boldsymbol{\lambda} = (\lambda_k, k = 0, \dots, N)$ is the vector of positive Lagrange multipliers associated with constraints (7) and (8). According to convex optimization theory, the primal-optimal vector \mathbf{X}^* and dual-optimal vector $\boldsymbol{\lambda}^*$ must satisfy Karush-Kuhn-Tucker (KKT) conditions [44]:

$$\begin{aligned} \nabla_{\mathbf{X}} L(\mathbf{X}^*, \boldsymbol{\lambda}^*) &= \mathbf{0}, \quad \boldsymbol{\lambda}^* \geq \mathbf{0}, \\ \sum_{n=1}^N a_{nk} x_{nk}^* &\leq r_k + \frac{M_k}{\delta}; \quad k = 1, \dots, N, \end{aligned} \quad (10)$$

$$\sum_{n=1}^N \sum_{k=1}^N a_{nk} x_{nk}^* \leq CR, \quad (11)$$

$$\lambda_0^* \left(\sum_{n=1}^N \sum_{k=1}^N a_{nk} x_{nk}^* - CR \right) = 0, \quad (12)$$

$$\lambda_k^* \left(\sum_{n=1}^N a_{nk} x_{nk}^* - r_k - \frac{M_k}{\delta} \right) = 0; \quad k = 1, \dots, N. \quad (13)$$

Expanding the first KKT condition, we have for any k and n :

$$\frac{\partial L}{\partial x_{nk}^*} = a_{nk} U'_{nk}(x_{nk}^*) - a_{nk} \lambda_k^* - a_{nk} \lambda_0^* = 0, \quad (14)$$

which further gives

$$x_{nk}^* = \left[a_{nk} U_{nk}'^{-1} (\lambda_0^* + \lambda_k^*) \right]^+, \quad (15)$$

where $[z]^+ = \max\{z, 0\}$. Having computed \mathbf{X}^* , we define the dual function $D(\cdot)$ as [44]:

$$D(\lambda) := \max_{\mathbf{X}} L(\mathbf{X}, \lambda) = L(\mathbf{X}^*, \lambda). \quad (16)$$

We note that Lagrange multipliers are referred to as dual variables since the dual function $D(\cdot)$ is a function of Lagrange multipliers. The dual-optimal vector λ^* could be found by solving the dual problem of problem (6)-(8), given by [44]:

$$\min_{\lambda \geq \mathbf{0}} D(\lambda). \quad (17)$$

Next, we solve dual problem (17) using iterative methods. Due to strict convexity of the primal problem (6)-(8), $D(\cdot)$ is continuously differentiable with partial derivatives given by Danskin's Theorem [43]:

$$\frac{\partial D}{\partial \lambda_k} = r_k + \frac{M_k}{\delta} - \sum_{n=1}^N a_{nk} x_{nk}; \quad k = 1, \dots, N, \quad (18)$$

$$\frac{\partial D}{\partial \lambda_0} = CR - \sum_{n=1}^N \sum_{k=1}^N a_{nk} x_{nk}. \quad (19)$$

Thus, we take the advantage of using simple iterative methods such as gradient projection algorithm and its variants. In order to achieve a fast yet simple iterative algorithm, we use *diagonally scaled gradient projection algorithm* [43] to solve problem (17). This algorithm can be seen as an approximate to the Newton's method and hence we expect that it will converge faster than the gradient method while exhibiting tractable scalability features.

Using this algorithm, the dual variable update equation at m -th iteration is given by

$$\lambda^{(m+1)} = \left[\lambda^{(m)} - \gamma^{(m)} \mathbf{B}^{(m)} \nabla D(\lambda^{(m)}) \right]^+, \quad (20)$$

where $\gamma^{(m)}$ is a step size and $\mathbf{B}^{(m)}$ is a diagonal matrix whose main diagonal elements are the inverse of second partial derivatives of the dual function:

$$\mathbf{B} = \text{diag} \left(\left(\frac{\partial^2 D}{\partial \lambda_k^2} \right)^{-1}, k = 0, \dots, N \right). \quad (21)$$

Using (18)-(19), the required second derivatives of the dual function stated in (16) are obtained below:

$$\frac{\partial^2 D}{\partial \lambda_k^2} = - \sum_{n=1}^N a_{nk} \frac{\partial x_{nk}}{\partial \lambda_k} = - \sum_{n=1}^N a_{nk} \frac{\partial}{\partial \lambda_k} U'_{nk}^{-1}(\lambda_0 + \lambda_k) = - \sum_{n=1}^N \frac{a_{nk}}{U''_{nk}(x_{nk})}, \quad k = 1, \dots, N, \quad (22)$$

$$\frac{\partial^2 D}{\partial \lambda_0^2} = - \sum_{n=1}^N \sum_{k=1}^N a_{nk} \frac{\partial x_{nk}}{\partial \lambda_0} = - \sum_{n=1}^N \sum_{k=1}^N a_{nk} \frac{\partial}{\partial \lambda_k} U'_{nk}^{-1}(\lambda_0 + \lambda_k) = - \sum_{n=1}^N \sum_{k=1}^N \frac{a_{nk}}{U''_{nk}(x_{nk})}. \quad (23)$$

Step size $\gamma^{(m)}$ must be chosen so as to guarantee the convergence of the iterative algorithm while yielding fast convergence. One of the best choices of step size is the *diminishing step size rule* that satisfies [43]:

$$\gamma^{(m)} \geq 0, \quad \forall m \in \mathbb{N}, \quad \lim_{m \rightarrow \infty} \gamma^{(m)} = 0, \quad \sum_{m=1}^{\infty} \gamma^{(m)} = \infty.$$

In this paper we use $\gamma^{(m)} = \frac{d}{\sqrt{m}}$ with some constant $d > 0$ (to chosen later), which satisfies the above conditions. Substituting partial derivatives into (20), update equations for dual variables λ_0 to λ_N are given by:

$$\lambda_0^{(m+1)} = \left[\lambda_0^{(m)} + \frac{\gamma^{(m)}}{\sum_{n=1}^N \sum_{k=1}^N \frac{a_{nk}}{U''_{nk}(x_{nk}^{(m)})}} \left(CR - \sum_{n=1}^N \sum_{k=1}^N a_{nk} x_{nk}^{(m)} \right) \right]^+, \quad (24)$$

$$\lambda_k^{(m+1)} = \left[\lambda_k^{(m)} + \frac{\gamma^{(m)}}{\sum_{n=1}^N \frac{a_{nk}}{U''_{nk}(x_{nk}^{(m)})}} \left(r_k + \frac{M_k}{\delta} - \sum_{n=1}^N a_{nk} x_{nk}^{(m)} \right) \right]^+, \quad k = 1, \dots, N, \quad (25)$$

where $x_{nk}^{(m)}$ is the value for the primal variables at the m -th iteration of the algorithm given by

$$\mathbf{X}^{(m+1)} = \arg \max_{\mathbf{X}} L(\mathbf{X}, \lambda^{(m)}). \quad (26)$$

Equivalently, $x_{nk}^{(m)}$ is given by (15) when $\lambda^{(m)}$ is used instead λ^* as its approximation until the m -th iteration.

Using the iterations outlined in (24) and (25), the sequence $\{\lambda^{(m)}\}$ will converge to dual-optimal variables λ^* . Because of strict convexity of primal problem (i.e., (6)-(8)), strong duality holds and hence it is guaranteed that both dual and primal problems will have the same optimal objective [44]. Thus, solving the dual problem (17) leads to the optimal solution of the primal and guarantees that the sequence $\{\mathbf{X}^{(m)}\}$ obtained by (26) converges to the optimal solution of the admission control problem. We defer the algorithmic aspects of this iterative solution until the next section.

6.2. α -Fair Utility Functions

Now we concentrate on the case of α -fair utility functions. First, we consider the case of $\alpha \neq 1$. Denoting by $\mathbf{X}^*(\alpha)$ the optimal rate vector for utility with parameter α , using (4) and (15) we have

$$x_{nk}^*(\alpha) = a_{nk} w_{nk} (\lambda_k + \lambda_0)^{-\frac{1}{\alpha}}. \quad (27)$$

Furthermore, calculation of partial derivatives for α -fair utility functions produces update equations for this class of utility functions given by (24) and (25) with $U''(x_{nk}^{(m)}) = w_{nk}[x_{nk}^{(m)}]^{-(\alpha+1)}$. It is worth noting that 1-fair utility ($\alpha = 1$) can be obtained by calculating the limit of α -fair utility for $\alpha \neq 1$ when α approaches 1:

$$x_{nk}^*(1) = \lim_{\alpha \rightarrow 1} x_{nk}^*(\alpha) = \frac{a_{nk}w_{nk}}{\lambda_k + \lambda_0}. \quad (28)$$

Accordingly, dual variable updates for $\alpha = 1$ can be obtained by asserting $\alpha = 1$ into (24) and (25). Thus, when describing the admission control algorithm in later sections, equations (27)-(25) are valid for all $\alpha > 0$.

7. Admission Control Algorithm

Although the above-mentioned optimization procedure for the admission control might look complicated, it could practically be implemented using the algorithms that will be described in this section. To gain more insights into the design of algorithms, we would only concentrate on the case of α -fair utility functions. We note, however, that one can simply use (15), (24), and (25) when using any other utility functions satisfying conditions **C1** and **C2**. In the sequel, we first present an algorithm for the iterative calculation of primal-optimal and dual-optimal vectors. Second, based on the obtained optimal values, we devise our final wavelength assignment algorithm by means of a centralized admission controller.

7.1. Iterative Solution to Admission Control Problem

Algorithm 1 lists the required steps to solve our admission control optimization problem (6)-(8) by iteratively solving its dual problem (17). This algorithm includes the iterations given by (27), (24), and (25), and begins by choosing an initial feasible value for \mathbf{X} and λ . As long as the specified stopping criterion (defined below) is not met, at each iteration, the dual variable vector $\lambda^{(m)}$ is updated first. Based on the updated dual variables, the algorithm calculates $\mathbf{X}^{(m)}$. Eventually the stopping criterion is met, and then $\mathbf{X}^{(m)}$ and $\lambda^{(m)}$ are reported as the approximate values of \mathbf{X}^* and λ^* , respectively.

One interesting performance metric is the convergence behavior of the proposed algorithm. Of particular concern is the number of iterations required for the algorithm to converge. Gradient-like algorithms do not converge after a finite number of iterations [43], and hence it is necessary to determine a stopping criterion for the algorithm based on some predefined accuracy that can approximate the distance from the globally optimal point.

Algorithm 1 Iterative Solution

Initialization:

Choose feasible starting points $\mathbf{X}^{(0)}$ and $\lambda^{(0)}$. Set $m = 1$.

while $\max_{n,k} |x_{nk}^{(m+1)} - x_{nk}^{(m)}| > \epsilon R$ **do**

Set step size $\gamma^{(m)} = \frac{d}{\sqrt{m}}$.

Update dual variables using (24) and (25).

Update primal variables for all k and n : $x_{nk}^{(m+1)} = \left[a_{nk} w_{nk} \left(\lambda_k^{(m+1)} + \lambda_0^{(m+1)} \right)^{-\frac{1}{\alpha}} \right]^+$.

Increment m : $m = m + 1$.

end while

The stopping criterion for the iterative algorithm is chosen as follows. Given $\epsilon \in (0, 1)$, we terminate the algorithm when the largest change in source rates is less than ϵR , i.e.,

$$\max_{n,k} |x_{nk}^{(m+1)} - x_{nk}^{(m)}| \leq \epsilon R. \quad (29)$$

We then define $I_\epsilon \triangleq \min \left\{ m : \max_{n,k} |x_{nk}^{(m+1)} - x_{nk}^{(m)}| \leq \epsilon R \right\}$ to denote the minimum number of iterations required for the algorithm to meet the stopping criterion mentioned above. We note that there is a tradeoff in choosing ϵ . A lower value of ϵ guarantees that the final result will be closer to the globally optimal point, however, at the expense of larger I_ϵ , i.e., more iterations.

7.1.1. Numerical Experiments

In order to quantify I_ϵ , we need to gain some insight into how the problem parameters would affect I_ϵ . To this end, we have carried out several simulation experiments through implementing Algorithm 1 in MATLAB. In particular, we are interested in the influence of the following parameters on I_ϵ :

1. Number of nodes in the crossbar (N)
2. Density and pattern of traffic matrix (\mathbf{A})
3. Step size (γ)

In all experiments, we consider a MWMR crossbar consisting of 32 waveguides each carrying 64 wavelengths. Considering 10 Gbps transmission rate for each wavelength, the total capacity of the crossbar is $C = 20.48$ Tbps. Moreover, for all nodes, the values for the parameters r_k and w_{nk} are drawn independently and uniformly at random from the intervals $[0, C]$ and $[0, 1]$, respectively. Each node is assumed to have free buffer space for storing g packets, where g is chosen from the set $\{1, 2, \dots, 20\}$ uniformly at random. Finally, we set $\delta = 5.4$ ns and $\epsilon = 10^{-11}$.


Table 1: Variance of number of iterations I_ϵ for different values of N , different densities of matrix \mathbf{A} , and step size $\gamma = \frac{5}{\sqrt{m}}$

$N \backslash$ Density of \mathbf{A}	0.5%	2%	10%	50%	90%
64	44.35	7.21	5.61	3.19	3.63
128	8.49	3.17	4.93	1.97	1.91
256	2.18	0.60	1.47	0.50	0.82


We consider crossbars with $N = 64, 128$, and 256 number of nodes, and use a diminishing step size $\gamma = \frac{d}{\sqrt{m}}$ with $d = 3, 5$, and 7 . Since both pattern and density of the traffic matrix \mathbf{A} may influence I_ϵ , we use randomly generated traffic matrices with different densities, ranging from sparse to very dense. The matrix density represents the ratio of non-zero elements (equivalently 1s) to all N^2 elements of \mathbf{A} . Thus, for example 2% density means that $0.02N^2$ elements of \mathbf{A} are 1. To take into account the randomness of problem parameters, for each choice of (N, \mathbf{A}, γ) , we report the empirical average of I_ϵ , denoted by \bar{I}_ϵ , measured over 100 experiments.

Figure 3 portrays \bar{I}_ϵ along with its 90% confidence interval for $N = 64, 128$, and 256 . We see that even with increase in the size of the crossbar N , the quantity \bar{I}_ϵ increases intangibly. This confirms that Algorithm 1 exhibits very good scalability properties, making it useful for many-core systems. Such a tractable scalability stems mainly from salient scalability properties of the diagonally scaled gradient projection algorithm applied in Section 6 to solve the dual problem. The results also imply that depending on N and the density of \mathbf{A} , each of the three chosen step sizes might be preferable. For example, for dense \mathbf{A} , the step size $\gamma = \frac{5}{\sqrt{m}}$ would require fewer iterations. Considering all of the above cases reveals that \bar{I}_ϵ diminishes at most to 50% of the choice with the maximum magnitude. This might sound a dramatic decrease in the average number of required iterations \bar{I}_ϵ . However, for most cases, such a reduction will be at most about 20-25 iterations. Table 1 lists the variance of I_ϵ for several combinations of traffic matrix densities and crossbar sizes with step size $\gamma = \frac{5}{\sqrt{m}}$.


Figure 3 along with Table 1 show that confidence intervals gradually shrink as the density of the traffic matrix \mathbf{A} increases. This is because increasing the density of \mathbf{A} is equivalent to decreasing the randomness in the pattern of the matrix. Thus, many experiments will encounter (partially) similar traffic patterns for which the I_ϵ values will be very close. This results in a decrease in the variance of I_ϵ , and hence, causes the confidence interval of \bar{I}_ϵ to shrink.



(a) Crossbar with $N = 64$ nodes



(b) Crossbar with $N = 128$ nodes



(c) Crossbar with $N = 256$ nodes

Figure 3: Average number of iterations \bar{I}_ϵ for a crossbar with three different number of nodes and $\gamma = \frac{d}{\sqrt{m}}$ for $d = 3, 5, 7$.

To summarize, the experiments reported above corroborate that (i) Algorithm 1 possesses very good and tractable scalability properties, and (ii) it requires only a few tens of iterations (about 25-50) to achieve a solution with high accuracy for a wide range of traffic patterns, choices of step sizes, and number of nodes.

7.2. Trimming Optimal Rates

In an optical on-chip crossbar, the rate granularity of each data channel is equal to the corresponding rate of one single wavelength R . As a result, allocated rates in practical scenarios must be a multiple of R . The optimal rates calculated by Algorithm 1, i.e., \mathbf{X}^* , might not satisfy this property and therefore, some rounding might be necessary. Let $\hat{\mathbf{X}}$ denote the optimal rate vector after the rounding process. In order to compute \hat{x}_{nk} from x_{nk}^* , we first decrement x_{nk}^* to the nearest legal value, and then define S as

$$S = \sum_{n=1}^N \sum_{k=1}^N (x_{nk}^* \bmod R).$$

To derive $\hat{\mathbf{X}}$, we propose a procedure outlined in Algorithm 2. This algorithm consists in choosing $\hat{x}_{nk} \in \hat{\mathbf{X}}$ at random with probability proportional to $x_{nk}^* - \hat{x}_{nk}$, and then incrementing \hat{x}_{nk} by R (provided that feasibility conditions are preserved) and decrementing S by R .

Algorithm 2 Trimming Optimal Rates

Calculate for all k and n : $\hat{x}_{nk} = x_{nk}^* - (x_{nk}^* \bmod R)$.

Calculate $S = \sum_{n=1}^N \sum_{k=1}^N (x_{nk}^* \bmod R)$.

while $S > 0$ **do**

Let $\mathcal{X} = \{\hat{x}_{nk} : x_{nk}^* - \hat{x}_{nk} > 0, \sum_{n=1}^N \hat{x}_{nk} + R \leq \frac{M_k}{\delta} + r_k, \forall n, \forall k\}$.

Choose an element of \mathcal{X} randomly with probability proportional to $x_{nk}^* - \hat{x}_{nk}$ and set: $\hat{x}_{nk} = \hat{x}_{nk} + R$.

$S \leftarrow S - R$

end while

7.3. The Case of Bursty Traffic

We propose another solution procedure which, in contrast to Algorithm 1, is not iterative. Yet it outputs a rate allocation which is optimal in some cases. Such a non-iterative procedure is very fast and efficient, and thus may prove useful for the case of bursty traffic. This procedure, described in Algorithm 3, is motivated as follows. First note that Algorithm 3 always outputs a feasible point for problem (6). Namely, its output satisfies constraints (7)-(8). Moreover, recall from (27) that for the case of proportionally fair utility functions

(i.e., $\alpha = 1$), optimal rates can be given by

$$x_{nk}^* = \frac{a_{nk}w_{nk}}{\lambda_k^* + \lambda_0^*}. \quad (30)$$

In order to justify that the output of Algorithm 3 is a good approximate solution, we consider two cases.

Case 1: Under-utilized regime. In this case, we assume that under optimal rate allocation, the whole resources of the system are not used by the nodes. Namely, we would have: $\sum_{n=1}^N \sum_{k=1}^N x_{nk}^* < CR$. Hence, KKT condition (12) implies that $\lambda_0^* = 0$, and therefore (30) gives

$$x_{nk}^* = \frac{a_{nk}w_{nk}}{\lambda_k^*}, \quad \forall k, \forall n. \quad (31)$$

Hence, we obtain

$$\frac{a_{n'k}w_{n'k}}{x_{n'k}^*} = \frac{a_{nk}w_{nk}}{x_{nk}^*} = \lambda_k^*, \quad \forall n, n'.$$

Now consider receiver node k . First note that in this case $\lambda_k^* > 0$ and hence, KKT condition (13) implies that

$$\sum_{n=1}^N a_{nk}x_{nk}^* = r_k + \frac{M_k}{\delta}.$$

Combing these two last relations, we obtain

$$x_{nk}^* = \frac{a_{nk}w_{nk}}{\sum_{j=1}^N a_{jk}w_{jk}} \left(r_k + \frac{M_k}{\delta} \right), \quad \forall n. \quad (32)$$

Finally, since the system is under-utilized, we will have $S = CR - \sum_{n=1}^N \sum_{k=1}^N x_{nk}^* > 0$, and therefore the output of Algorithm 3 is given by (32). Thus, Algorithm 3 gives the optimal solution for the under-utilized case.

We remark that the solution (32) is intuitive: for any receiver node k , the quantity $\sum_{j=1}^N a_{jk}w_{jk}$ is the sum of the weights of the nodes who wish to send packets to node k . Thus, (32) implies that the optimal allocation shares the available capacity $r_k + \frac{M_k}{\delta}$ among nodes $(n, a_{nk} = 1)$ proportionately to their weights.

Case 2: Fully-utilized regime. In this case, at the optimal point we will have $\sum_{n=1}^N \sum_{k=1}^N x_{nk}^* = CR$, and so by KKT condition (12), we will have $\lambda_0^* > 0$. It then follows from (30) that

$$x_{nk}^* = \frac{a_{nk}w_{nk}}{\lambda_k^* + \lambda_0^*} < \frac{a_{nk}w_{nk}}{\lambda_k^*}, \quad \forall k, \forall n. \quad (33)$$

Hence, for any n and k , if we use the rate given by (32), i.e., choose

$$x_{nk} = \frac{a_{nk}w_{nk}}{\sum_{j=1}^N a_{jk}w_{jk}} \left(r_k + \frac{M_k}{\delta} \right), \quad \forall n, \quad (34)$$

the constraint (8) will be satisfied. However, we have that $x_{nk}^* < x_{nk}$, and since the system is fully-utilized, we will have

$$CR - \sum_{n=1}^N \sum_{k=1}^N a_{nk} x_{nk} < CR - \sum_{n=1}^N \sum_{k=1}^N a_{nk} x_{nk}^* = 0.$$

Thus, constraint (7) will be violated (equivalently $S < 0$). To resolve this issue, we consider the uniform rate allocation $\frac{CR}{\sum_n \sum_k a_{nk} x_{nk}}$ and set x_{nk} to the minimal value between (34) and the uniform allocation. It then follows that this latter choice of x_{nk} , $\forall n, k$ satisfies both constraints. Thus, in this case, Algorithm 3 generates a sub-optimal but feasible rate allocation. Finally, we remark that when $\lambda_0^* \ll \lambda_k^*$, $\forall k$, this solution is near-optimal as verified by (33).

Algorithm 3 Solution for Burst Traffic

Compute for all n and k : $z_{nk} = \frac{a_{nk} w_{nk}}{\sum_{j=1}^N a_{jk} w_{jk}} \left(r_k + \frac{M_k}{\delta} \right).$

Compute: $S = CR - \sum_{n=1}^N \sum_{k=1}^N a_{nk} z_{nk}.$

if $S \geq 0$ **then**

Set $x_{nk}^* = z_{nk}$ for all n, k .

else

Set for all n and k : $x_{nk}^* = \min \left(z_{nk}, \frac{CR}{\sum_{n=1}^N \sum_{k=1}^N a_{nk}} \right).$

end if

7.4. α -Fair Admission Control Algorithm

We now describe the algorithm for optimal wavelength assignment and buffer management based on a central on-chip admission controller. The algorithm is listed below as Algorithm 4. A built-in admission controller, which from now on we refer to as the controller, is supposed to be mounted onto the system to implement this algorithm. At each time slot t , all nodes are required to send their requests to the controller. The request of each node consists of the set of its target destinations, and corresponds to a row in the matrix \mathbf{A} . Moreover, the request contains information about each node's available buffer space as well as the rate at which incoming packets can be processed. After receiving all requests, the controller acquires a_{nk} , w_{nk} , M_k , and r_k for every n and k . Then, the controller calculates the optimal rate allocation using Algorithm 1 or 3. According to the final obtained values for rate allocations $\hat{\mathbf{X}}$, the controller assigns distinct wavelengths on the available waveguides to each node for sending and receiving data.

From computational complexity point of view, the proposed admission controller must be capable of doing simple mathematical and logical operations to implement Algorithm 3. It is worth noting that the

Algorithm 4 α -Fair Admission Control Algorithm for Buffer Management and Wavelength Assignment

At time slot t :

Get M_k^{t+1} , r_k^{t+1} , a_{nk}^{t+1} and w_{nk}^{t+1} for $n, k = 1, \dots, N$.

Calculate optimal rate values for time slot $t + 1$ using Algorithm 1 or 3.

Trim the values acquired in Step 2 using Algorithm 2.

Based on the rate allocation values from Step 3, inform each node of the channels it should use in time slot $t + 1$.

output of the controller should be communicated to nodes in a simple yet fast manner. This may require the system to be equipped with auxiliary optical/electrical connections, as a dedicated signaling media, to communicate admission control results from the controller to nodes as well as delivering requests in the reverse direction. Such a separate media that decouples transmission of data packets and control packets calls for the recently appreciated SDN architectures in networking research community that decouple network control and forwarding functions. This architecture has also been employed in a nanophotonic NoC named 2D-HERT [45]. At each time slot, nodes start sending and receiving data over the crossbar based on the results received from the controller. Meanwhile, nodes send to the controller their requests for the next time slot. Thus, data and control communications can be overlapped so that admission control results are readily accessible by nodes at the beginning of each time slot and controller-related overheads are removed. A similar approach has also been utilized in [46] to increase the channel efficiency of a wireless NoC MAC protocol.

It is worth noting that the overlap between data and control packets can be achieved only if the length of each time slot δ is greater than or equal to the time it takes to send the requests to the controller and get the results back. This condition determines the minimum required bandwidth and the number of waveguides for communications to/from the controller. In this regard, let $rqst$ and $rslt$ respectively represent the size of the request and result vectors communicated between each node and the controller. Moreover, let B_c denote the bandwidth of the channel between each node and the controller. Then, the required time for sending the requests to the controller and getting the results back would be at most $\frac{rqst}{B_c} + \frac{rslt}{B_c} + V$, where V is the maximal overhead due to controller's computations. Thus, we require

$$\frac{rqst}{B_c} + \frac{rslt}{B_c} + V \leq \delta, \quad (35)$$

or equivalently: $B_c \geq \frac{rqst+rslt}{\delta-V}$.

From (35), we can also compute the total number of waveguides required to realize the communication channels between the central controller and the on-chip nodes. For instance, in a crossbar with N nodes and

64 wavelengths with rate R multiplexed on each waveguide, the number of waveguides W_c should satisfy

$$W_c \geq \lceil \frac{NB_c}{64R} \rceil \geq \lceil \frac{N(\text{rqst} + \text{rslt})}{64R(\delta - V)} \rceil. \quad (36)$$

8. Hardware Implementation

In this section a practical implementation of the proposed controller is described. It is followed by an analysis of its area and power consumption overheads.

8.1. Hardware Design

In order to account for both power and latency limitations as the key concerns for the design of nowadays digital systems, we propose two variants: A *low latency* implementation that strives to output the results with minimal latency, and a *low power* implementation whose aim is consume the least possible power. Figure 4 portrays the basic structure of the proposed implementation for both variants. To facilitate the presentation, the concentration in the subsequent design will be on implementation of an iterate of Algorithm 1, thereby ignoring Algorithm 2⁴. To alleviate implementation complexity, we use fixed point representation for real numbers. In Algorithm 1, each iteration m consists of three basic steps:

- (i) Setting step size $\gamma^{(m)}$
- (ii) Updating dual variables $\lambda_0^{(m+1)}, \lambda_k^{(m+1)}$, for all $k = 1, \dots, N$
- (iii) Calculating primal variables $x_{nk}^{(m+1)}$ for all n, k such that $a_{nk} = 1$

Next we describe how to implement these three steps in detail.

Step (i). This step, which has the same implementation for both variants, is implemented using a look-up table. Namely, we calculate all values of step size at design time and store them in look-up table array.

Step (ii). Computation of this step is divided into three phases denoted by F1, F2, and F3 in Figure 4. In phase F1, to achieve a high performance controller for low latency implementation, the terms $\sum_{n=1}^N \frac{a_{nk}}{w_{nk}} [x_{nk}^{(m)}]^2$ and $\sum_{n=1}^N a_{nk} x_{nk}^{(m)}$ for $k = 1, \dots, N$ are computed in a fully parallel structure. To this end, we first calculate $\frac{a_{nk}}{w_{nk}} [x_{nk}^{(m)}]^2$ and $a_{nk} x_{nk}$ for all n, k , and then accumulate the calculated numbers of each column.

⁴The overhead due to Algorithm 2 is ignored as an alternative of this algorithm can be simply implemented using shift operation and XOR-based pseudo-random number generation. Furthermore, for practical purposes we propose to run Algorithm 1 for a fixed number of iterations, in contrast to the conditional stopping.

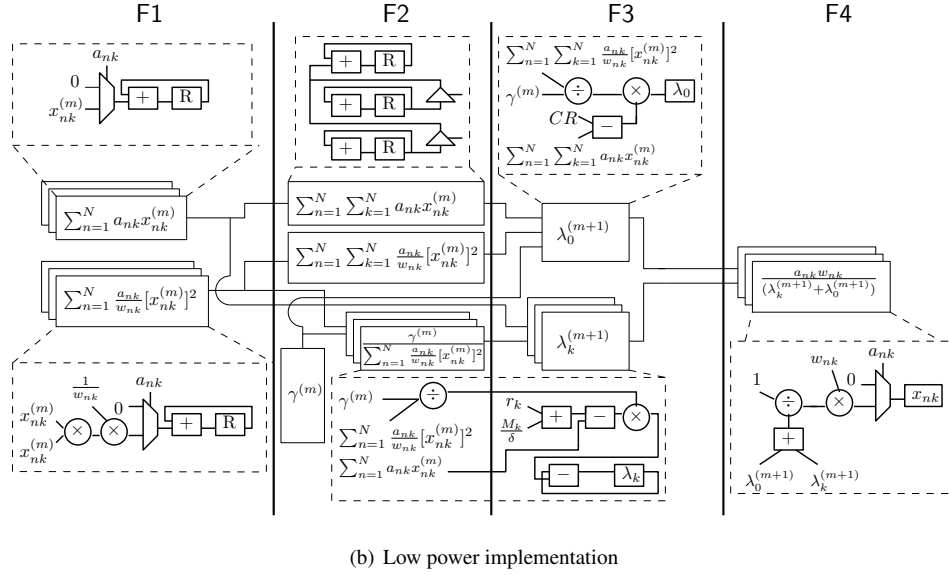
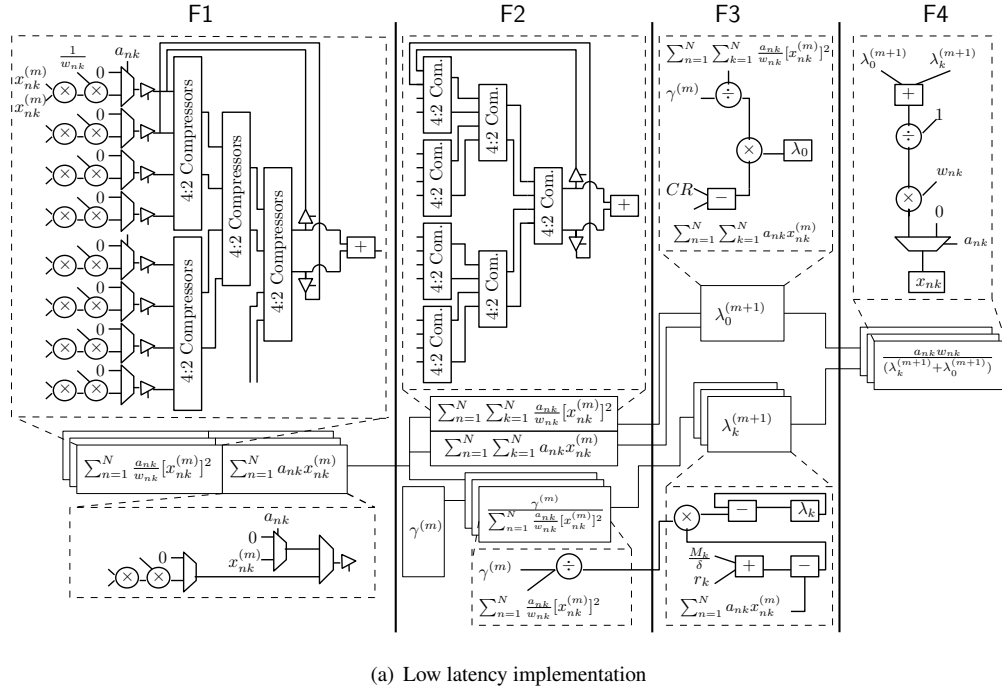


Figure 4: The basic structure of different implementations of the proposed controller.

Recall that $a_{nk} \in \{0, 1\}$ and we simply calculate $a_{nk}x_{nk}^{(m)}$ by a multiplexer. Moreover, to reduce computation complexity of this phase, we store both of $\frac{1}{w_{nk}}$ and w_{nk} ⁵. Thus, we can compute $\frac{1}{w_{nk}}[x_{nk}^{(m)}]^2$ by using one multiplier and two multiplexers, and control the multiplier inputs by the controller part of this circuit. Finally, for each column k , we accumulate different numbers $\frac{a_{nk}}{w_{nk}}[x_{nk}^{(m)}]^2$ for $n = 1, \dots, N$ in parallel. To attain high performance circuit for this accumulation, we use a combination of carry save idea [47] and tree structure summation as shown in Figure 4. Moreover, sharing this circuit for computations of $\sum_{n=1}^N \frac{a_{nk}}{w_{nk}}[x_{nk}^{(m)}]^2$ and $\sum_{n=1}^N a_{nk}x_{nk}^{(m)}$, the hardware overhead of the proposed controller will be significantly reduced. To this end, the proposed implementation computes summation $\sum_{n=1}^N a_{nk}x_{nk}^{(m)}$ when it calculates $\frac{a_{nk}}{w_{nk}}[x_{nk}^{(m)}]^2$.

Unlike the low latency variant, in phase F1 of the low power variant, the terms $\sum_{n=1}^N a_{nk}x_{nk}^{(m)}$ and $\sum_{n=1}^N \frac{a_{nk}}{w_{nk}}[x_{nk}^{(m)}]^2$ for $k = 1, \dots, N$ are computed in partially parallel manner. As a tradeoff between achieving higher performance and lower area and power overhead, we use a pipeline structure similar to multiple and accumulate operation in phase F1. In other words, implementation of this variant performs computation of $\frac{1}{w_{nk}}[x_{nk}^{(m)}]^2$ and the summation operation simultaneously. Furthermore, to obtain acceptable performance in low power variant we compute $\sum_{n=1}^N \frac{a_{nk}}{w_{nk}}[x_{nk}^{(m)}]^2$ and $\sum_{n=1}^N a_{nk}x_{nk}^{(m)}$ in parallel using distinct circuits.

Two types of computations are accomplished in phase F2: summation ($\sum_{n=1}^N \sum_{k=1}^N a_{nk}x_{nk}^{(m)}$, $\sum_{n=1}^N \sum_{k=1}^N \frac{a_{nk}}{w_{nk}}[x_{nk}^{(m)}]^2$) and division ($\frac{\gamma^{(m)}}{\sum_{n=1}^N \frac{a_{nk}}{w_{nk}}[x_{nk}^{(m)}]^2}$). For summation, results of phase F1 are accumulated. To mitigate hardware overhead of low latency implementation, we use the summation circuit used in phase F1 to accumulate N different numbers. However, in low power implementation to attain high performance controller, we use tree structure summation to accumulate N different numbers. Moreover, to mitigate hardware overhead of low power variant, we share the required adders of this computation with adders of phase F1 of this variant. For division, computation of all columns is performed simultaneously for both variants. As in [48], we use multiplier based divider, which is popular in commercial processors. Hence, sharing required multipliers of division with multipliers of phase F1 further reduces hardware overhead. Finally, the circuit of phase F3 performs last computations of updating dual variables in parallel for both implementations. Again, to alleviate the hardware cost, we share the computational resources of this phase with the computational resources of phase F1.

Step (iii). Last phase of each iterate in Algorithm 1 is devoted to the calculation of primal variables using dual variables. To this end, we first calculate the denominator of primal variables using dual variables and then compute all primal variables in a fully (resp. partially) parallel way for low latency (resp. low power)

⁵Indeed we require each node n to report both w_{nk} and $\frac{1}{w_{nk}}$, for $k = 1, \dots, N$.

variant. The basic structure of implementation of this phase is presented as phase F4 in Figure 4. To obtain a implementation with lower hardware overhead in low latency variant, we reuse multipliers of phase F1 in the division and multiplication computation of this phase. It is worth noting that by performing proper resource sharing among different phases in both implementations, the controller part of the implementation circuit controls the inputs of shared resources in datapath using multiplexers and selects correct inputs based on its state. As a tradeoff between achieving higher performance and lower hardware overhead for low power implementation, we calculate the elements of all columns of primal variables in parallel so that elements of each columns are computed using a pipeline approach.

8.2. Area Overhead

To have an intuition for the area overhead of the proposed controller, we consider a system with 64 cores, i.e., $N = 64$ in 32 nm process. To calculate storage overhead, we use CACTI [49]. Moreover, we estimate the computational overhead using the data provided [50] and [51] for computational operations. In this system, the total area overhead of the proposed implementation are 9.416539 mm^2 and 1.087179 mm^2 for low latency and low power implementations, respectively. In particular, our calculations show that only around $9.326724/662 = 1.5\%$ of the total chip area (e.g., in the case of Xeon E5-2699 v3) for low latency implementation will be occupied by the controller. Similarly, the area overhead for low power implementation is nearly $1.087179/662 = 0.2\%$ of the total chip area, signifying that the area overhead of the proposed controller in the case of both variants is negligible.

8.3. Power Consumption Overhead

We now compute the power overhead of the controller. Unless stated otherwise, we consider a 32 nm process.

Dynamic power: Similarly to [52], to estimate the dynamic power consumption of the proposed controller, we use the method presented in McPAT tool [50]. To this end, we count the number of operations involved in each component of the presented implementation. Then, by multiplying this number by the energy consumption of each component that is consumed for each access, we obtain its dynamic energy consumption. Note that we adopt the dynamic energy consumption of computational component (such as multiplier, adder, etc.) and storage component using [51] and CACTI [49], respectively. Let $N = 64$ and assume that each iteration lasts for 27 (resp. 150) clocks for low latency (resp. low power) implementation. Hence, the total

dynamic power consumption of the proposed controller is 6.99 W and 1.45 W for low latency and low power variants, respectively.

Leakage Power: Using CACTI, we estimate the leakage power of storage parts of the proposed controller. Moreover, using McPAT together with data provided in [51], we calculate the leakage power of computational parts of the controller. We obtain that the total leakage power of the controller for low latency and low power variants are 4.96 W and 0.25 W, respectively. Finally, we find out that the power overhead of the proposed controller constitutes about $11.95/145 = 8.2\%$ and $1.7/145 = 1.2\%$ of the total power consumed by the chip (e.g., Xeon E5-2699 v3) for low latency and low power implementations, respectively. This highlights that the power overhead of the controller especially for low power implementation is negligible.

9. Simulation Experiments

For simulation experiments, we use OMNeT++ [53] to simulate an optical MWMR on-chip crossbar equipped with our proposed admission control policy. We will refer to such an architecture as MWMR-AC. We consider an MWMR-AC consisting of 64 nodes, each operating at 5 GHz clock speed. We also consider 64 waveguides for the data channels of the crossbar. Moreover, DWDM technique is used to carry 64 wavelengths on each waveguide simultaneously, providing a total of 4096 data channels. Each wavelength is considered to provide a bandwidth of $R = 10$ Gbps. Finally, in all the experiments, we focus on the case of $\alpha = 1$, which corresponds to *proportional fairness* metric [42].

In all experiments we assume $\delta = 6$ ns. This choice of δ follows from our estimate of controller’s minimal delay (5.4 ns) obtained from our hardware implementation in Section 8. Moreover, based on (36) and in accordance to our simulated crossbar, we consider 20 waveguides to realize the communication channels between the central controller and the on-chip nodes.

For the sake of comparison, we conduct the experiments with the Corona crossbar [34], with 64 crossbar waveguides each carrying 64 wavelengths simultaneously. Thus, each home node will have one dedicated waveguide (64 wavelengths) for data reception. We also consider Corona with the enhanced “Fast Forward” token arbitration mechanism (Corona-FF) proposed in [35].

We consider both synthetic and real traffic patterns in our simulation experiments. In particular, we use Uniform and Hot Spot for synthetic patterns, whereas SPLASH-2 and PARSEC are used for real benchmark evaluations.

9.1. Synthetic Patterns

In this subsection, we evaluate the performance of MWMR-AC for two synthetic patterns. In particular, we use the following three measures:

- (i) *Latency* defined as the difference between the time a packet arrives at the destination and the time it was generated. We focus on the average latency experienced by all delivered packets.
- (ii) *Network Throughput* that is the rate at which packets are delivered by the underlying network. We report the aggregate throughput normalized by the total crossbar capacity.
- (iii) *Nodal Throughput* defined as the rate at which each node can send packets to other destination nodes.


Network throughput is indeed the sum of all nodal throughputs across the crossbar.

In all cases, latency and throughput are measured as a function of the offered load injected to the network. Note that offered load values are normalized with respect to the total capacity provided by the crossbar. An offered load value of 1 represents the maximum load that can potentially be delivered by the crossbar.


Figure 5(a) and Figure 5(b) respectively show the latency and throughput results under the Uniform traffic pattern. Under the Uniform pattern, packet destinations are chosen randomly with a uniform distribution. As shown in Figure 5(a), MWMR-AC achieves lower latency than Corona/Corona-FF for offered loads above 0.2. We also see a mild increase in latency up to offered load 0.9 with MWMR-AC, whereas Corona's latency becomes unbounded above offered load 0.4. This is in line with the results shown for throughput in Figure 5(b). While MWMR-AC provides a maximum of 0.9 throughput, Corona/Corona-FF can only achieve a maximum throughput of 0.4. Moreover, MWMR-AC hits the saturation point at offered load 1.0, whereas Corona/Corona-FF is saturated at offered load 0.5. The saturation throughput itself (≈ 0.7) is also higher for MWMR-AC (0.7) compared to Corona/Corona-FF. We also see a higher saturated throughput for Corona-FF (0.4) compared to Corona (0.3). The Fast Forward token mechanism used in Corona-FF helps Corona to achieve a relatively higher throughput.

Another important observation from Figure 5(b) is the throughput decrease after offered load 0.9. This is the point where the controller starts using the iterative solution for channel assignments, i.e., it switches from Algorithm 3 to Algorithm 1. Consequently, the higher delay of the iterative solution causes about 0.2 drop in throughput. However, throughput is not decreased any further and is saturated around 0.7.

Latency results also show that for offered loads 0.1 and 0.2, Corona/Corona-FF provides slightly lower latency. This is due to the fact that with such low traffic loads, controller's delay becomes the dominant latency factor in MWMR-AC. In such cases, demand for communication resources is too low that prevents us from achieving any benefits from better sharing and efficient utilization of channels.




(a) Average latency; uniform traffic




(b) Network throughput; uniform traffic

Figure 5: Average latency and network throughput for the Uniform traffic pattern.



(a) Average latency; Hot Spot traffic



(b) Network throughput; Hot Spot traffic

Figure 6: Average latency and network throughput for the Hot Spot traffic pattern.

Figure 6 shows the latency and throughput results for the Hot Spot pattern. Under the Hot Spot pattern, all nodes send their packets to one single destination node. Without loss of generality, here we assume that node 0 is the Hot Spot destination. Latency results in Figure 6(a) show that MWMR-AC achieves a significantly lower latency compared to Corona/Corona-FF in the Hot Spot pattern. Note the different scale used in the two latency axes. In the case of MWMR-AC, the average latency is 3 orders of magnitude lower. This is because MWMR-AC has a much better sharing and utilization of all the channels that are provided by the crossbar. With the Hot Spot pattern, all nodes attempt to access the limited number of channels that are provided by the single waveguide that is dedicated to node 0 for data reception. The resulting high contention will also increase the overheads of the token-based arbitration mechanism used in Corona as token of node 0 will be without any credit most of the time. On the contrary, the admission control policy in MWMR-AC utilizes all the data channels available in the crossbar, and efficiently assigns them to the set of source nodes.

From Figure 6(b), we can also see that MWMR-AC achieves higher throughput compared to Corona/Corona-FF. While Corona/Corona-FF is already at the saturated point from the initial 0.01 offered load, we see linear increase in the throughput of MWMR-AC up to 0.13. After that, similarly to what we discussed for the case of Uniform pattern, the throughput is decreased to a saturated bound of 0.07. As shown, this is about 10 times higher than the saturation throughput of Corona/Corona-FF which is less than 0.007. It is worth mentioning that the maximum throughput achieved under the Hot Spot traffic is lower than that of Uniform traffic. This is totally expected; the limited buffer space and drain rate of the single destination node (i.e., node 0) limits the total achievable bandwidth for the Hot Spot pattern.

9.2. Fairness Analysis

In another experiment, we evaluate the behavior of our proposed admission control policy in terms of fairness. To this aim, we measure the nodal throughput experienced by each node for sending packets. This result highlights the portion of the capacity of crossbar allocated to each node. The Hot Spot traffic pattern with a high offered load value is used for these experiments since fairness issues mainly arise in the presence of contention for resources. We note, however, that our results with high-load Uniform traffic pattern (not shown here) exhibit a similar trend.

We first enumerate nodes from 0 to 63, and distribute them into three categories. In particular, nodes 0 to 20 belong to the first category, nodes with indices 21 to 41 belong to the second category, and the rest belong to the third category. Each category is then assigned with a weight so that all nodes in the same category will have the same weight for sending data to other receiving nodes.

We consider two scenarios in terms of the weights assigned to the three categories. In the first scenario, all categories will have the same weight (equal to 1), whereas in the second one, the weight for the first, the second, and the third category of nodes is respectively set to 1, 2, and 3.

Figure 7 illustrates the corresponding results. As shown in Figure 7(a), the nodal throughput seen by each node is proportional to its weight parameter. The bandwidth achieved by the nodes in the second and third categories is respectively two and three times higher than that of the nodes in the first category. Moreover, nodes belonging to the same category have achieved equal nodal throughput. Furthermore, as shown in Figure 7(b), the nodal throughput seen by all nodes are the same when all categories have the same value for the weight.

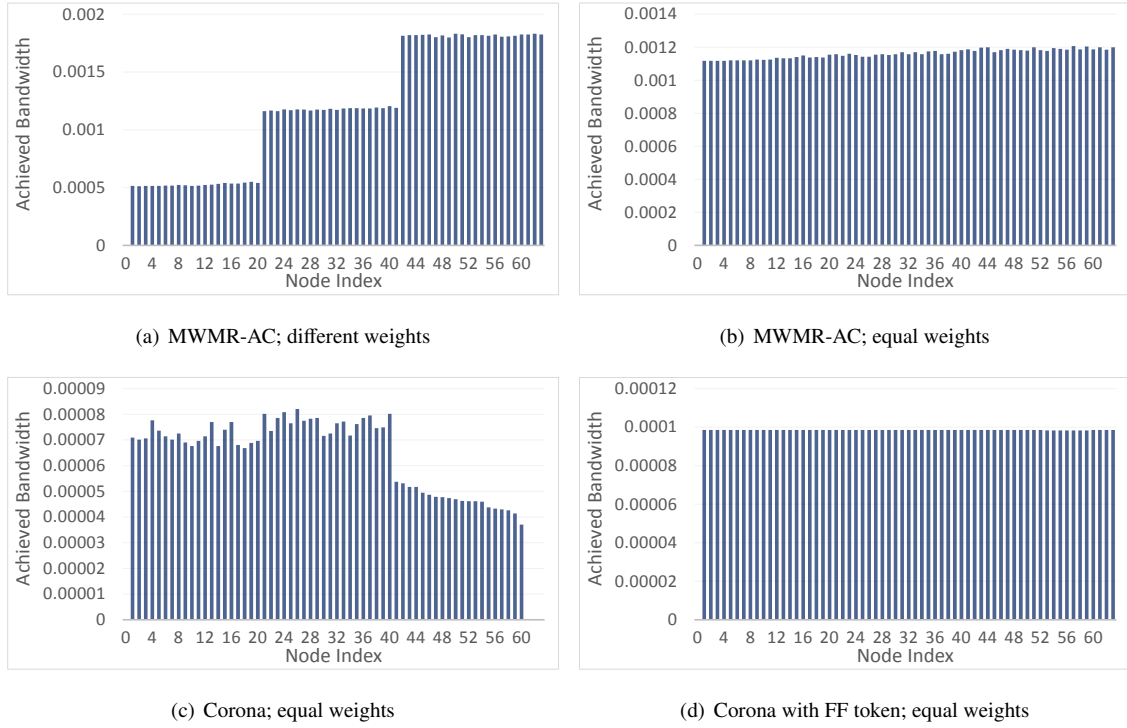


Figure 7: Nodal throughput of the nodes in MWMR-AC and Corona

For Corona/Corona-FF, we only have the results with all nodes having the same weight since Corona/Corona-FF does not provide any weighting or service differentiation capability. In fact, our admission control policy has the advantage of enabling service differentiation among nodes as shown in Figure 7(a). Figure 7(c) shows that Corona without Fast Forward token mechanism has an unfair nodal throughput across nodes.




Figure 8: Average latency of packets for various benchmarks from SPLASH-2 and PARSEC benchmark suites.

Nodes closer to node 0 achieve higher bandwidth than the farther ones. Specifically, nodes 61, 62, and 63 are totally starved. Corona-FF however, does provide a fair nodal throughput.

9.3. Parsec and SPLASH-2 Simulation Results

For the benchmark simulations, we use Sniper [54], an x86 multicore simulator, to generate PARSEC and SPLASH-2 traffic traces. The traces are then injected into our simulator to obtain the results. In Sniper, we consider 64 nodes, each having its own private L1/L2 (32kB/512kB) and shared L3 caches (1MB). We also consider a 1 (resp. 3) cycle latency to access the tag (resp. data) of L1 cache (4-way), a 3 (resp. 9) cycle latency to access the tag (resp. data) of L2 cache (8-way), a 4 (resp. 11) cycle latency to access the tag (resp. data) of L3 cache (16-way), a cache line size of 64 bytes, and a 100 ns latency to access the main memory. Furthermore, 8 memory controllers are used for accessing the main memory.

Figure 8 shows the total average latency of packets for each benchmark. It can be seen that MWMR-AC outperforms Corona and Corona-FF in 5 of the benchmarks, i.e., water.nsq, bodytrack, dedup, fluidanimate, and vips. In particular, for bodytrack and dedup, we see a significant decrease (4.7x and 6.04x) in latency compared to Corona and Corona-FF. For the rest of the benchmarks, both Corona and Corona-FF provide lower latency than MWMR-AC. This is mainly because these benchmarks impose lower traffic loads on the network in which case the latency of the controller in MWMR-AC becomes the bottleneck. This also

explains why we see very close latency results (about 10 to 11 ns) across these group of benchmarks for MWMR-AC. Note that essentially, a non-trivial admission control policy cannot provide much of a benefit compared to a trivial resource allocation mechanism when the traffic load is low. This is because with a low traffic, there is no contention for resources. In such cases, a fast and non-iterative solution such as the one provided in Algorithm 3 is more desirable. Although we use Algorithm 3 for low traffic loads, we have considered the overhead of this algorithm equal to one iteration in Algorithm 1. We believe that in practice, such overhead could be considerably lower for Algorithm 3, which will in turn improve the performance of MWMR-AC for low-traffic benchmarks.

10. Conclusion

Usage of WDM techniques in an optical crossbar provides a huge pool of wavelengths to be shared among competing on-chip cores. In order to manage such resources in a fair and efficient manner, we presented an admission control scenario in an optical on-chip Multiple Write, Multiple Read crossbar. In order to take into account the perceived satisfaction of cores when transmitting data, we cast the problem of wavelength assignment and buffer management as a utility-based convex optimization problem, also referred to as admission control problem. This formulation allowed us to devise a fair and efficient wavelength assignment and buffer management algorithm as the solution to admission control optimization problem. Running on a central admission controller, the proposed algorithm not only tries to achieve the maximum utilization of data channels, but also provides a mechanism to control the access of on-chip nodes to the shared communication resources.

References

- [1] G. E. Moore, Cramming more components onto integrated circuits, *Proceedings of the IEEE* 86 (1) (1998) 82–85.
- [2] D. Geer, Chip makers turn to multicore processors, *Computer* 38 (5) (2005) 11–13.
- [3] J. Parkhurst, J. Darringer, B. Grundmann, From single core to multi-core: preparing for a new exponential, in: *Proceedings of the 2006 IEEE/ACM international conference on Computer-aided design*, ACM, 2006, pp. 67–72.
- [4] ITRS, International technology roadmap for semiconductors (2009).

- [5] T. Bjerregaard, S. Mahadevan, A survey of research and practices of network-on-chip, *ACM Computing Surveys (CSUR)* 38 (1) (2006) 1–51.
- [6] W. J. Dally, B. Towles, Route packets, not wires: on-chip interconnection networks, in: *Proceedings of Design Automation Conference (DAC'01)*, IEEE, 2001, pp. 684–689.
- [7] L. Benini, G. De Micheli, Networks on chips: a new SoC paradigm, *Computer* 35 (1) (2002) 70–78.
- [8] D. Bertozzi, G. Dimitrakopoulos, J. Flich, S. Sonntag, The fast evolving landscape of on-chip communication, *Design Automation for Embedded Systems* (2014) 1–18.
- [9] J. D. Owens, W. J. Dally, R. Ho, D. J. Jayasimha, S. W. Keckler, L.-S. Peh, Research challenges for on-chip interconnection networks, *IEEE micro* 27 (5) (2007) 96–108.
- [10] R. A. Soref, J. P. Lorenzo, All-silicon active and passive guided-wave components for $\lambda = 1.3$ and 1.6μ , *IEEE Journal of Quantum Electronics* 22 (6) (1986) 873–879.
- [11] T. K. Woodward, A. V. Krishnamoorthy, 1-Gb/s integrated optical detectors and receivers in commercial CMOS technologies, *Selected Topics in Quantum Electronics*, *IEEE Journal of* 5 (2) (1999) 146–156.
- [12] V. R. Almeida, C. A. Barrios, R. R. Panepucci, M. Lipson, M. A. Foster, D. G. Ouzounov, A. L. Gaeta, All-optical switching on a silicon chip, *Optics Letters* 29 (24) (2004) 2867–2869.
- [13] C. Gunn, CMOS photonics for high-speed interconnects, *Micro*, *IEEE* 26 (2) (2006) 58–66.
- [14] D. A. B. Miller, Rationale and challenges for optical interconnects to electronic chips, *Proceedings of the IEEE* 88 (6) (2000) 728–749.
- [15] L. P. Carloni, P. Pande, Y. Xie, Networks-on-chip in emerging interconnect paradigms: Advantages and challenges, in: *Proceedings of the 2009 3rd ACM/IEEE International Symposium on Networks-on-Chip*, IEEE Computer Society, 2009, pp. 93–102.
- [16] A. Shacham, K. Bergman, L. P. Carloni, Photonic networks-on-chip for future generations of chip multiprocessors, *Computers*, *IEEE Transactions on* 57 (9) (2008) 1246–1260.
- [17] E. Bolotin, I. Cidon, R. Ginosar, A. Kolodny, QNoC: QoS architecture and design process for network on chip, *Journal of systems architecture* 50 (2) (2004) 105–128.

- [18] K. Goossens, J. Dielissen, J. van Meerbergen, P. Poplavko, A. Rădulescu, E. Rijpkema, E. Waterlander, P. Wielage, Guaranteeing the quality of services in networks on chip, in: *Networks on chip*, Springer, 2003, pp. 61–82.
- [19] B. Grot, S. W. Keckler, O. Mutlu, Preemptive virtual clock: a flexible, efficient, and cost-effective QoS scheme for networks-on-chip, in: *Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture*, ACM, 2009, pp. 268–279.
- [20] M. Chiang, S. H. Low, A. R. Calderbank, J. C. Doyle, Layering as optimization decomposition: A mathematical theory of network architectures, *Proceedings of the IEEE* 95 (1) (2007) 255–312.
- [21] C. A. D. Adi, H. Matsutani, M. Koibuchi, H. Irie, T. Miyoshi, T. Yoshinaga, An efficient path setup for a photonic network-on-chip, in: *Networking and Computing (ICNC), First International Conference on*, 2010, pp. 156–161. doi : 10.1109/IC-NC.2010.31.
- [22] M. Abdollahi, A. Namazi, S. Mohammadi, Clustering effects on the design of opto-electrical network-on-chip, in: *24th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP)*, 2016, pp. 427–430.
- [23] A. B. Ahmed, M. Meyer, Y. Okuyama, A. B. Abdallah, Hybrid photonic NoC based on non-blocking photonic switch and light-weight electronic router, in: *IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2015, pp. 56–61.
- [24] A. García-Guirado, R. Fernández-Pascual, J. M. García, S. Bartolini, Managing resources dynamically in hybrid photonic-electronic networks-on-chip, *Concurrency and Computation: Practice and Experience* 26 (15) (2014) 2530–2550.
- [25] G. Kurian, J. E. Miller, J. Psota, J. Eastep, J. Liu, J. Michel, L. C. Kimerling, A. Agarwal, Atac: a 1000-core cache-coherent processor with on-chip optical network, in: *Proceedings of the 19th international conference on Parallel architectures and compilation techniques*, ACM, 2010, pp. 477–488.
- [26] X. Wu, J. Xu, Y. Ye, Z. Wang, M. Nikdast, X. Wang, Suor: Sectioned unidirectional optical ring for chip multiprocessor, *ACM Journal on Emerging Technologies in Computing Systems (JETC)* 10 (4) (2014) 29.

- [27] M. Briere, B. Girodias, Y. Bouchebaba, G. Nicolescu, F. Mieyeville, F. Gaffiot, I. O'Connor, System level assessment of an optical noc in an MPSoC platform, in: *Proceedings of the conference on Design, automation and test in Europe*, 2007, pp. 1084–1089.
- [28] S. Koohi, S. Hessabi, Scalable architecture for a contention-free optical network on-chip, *Journal of Parallel and Distributed Computing* 72 (11) (2012) 1493–1506.
- [29] S. Koohi, Y. Yin, S. Hessabi, S. J. Yoo, Towards a scalable, low-power all-optical architecture for networks-on-chip, *ACM Transactions on Embedded Computing Systems (TECS)* 13 (3) (2014) 101:1–101:30.
- [30] S. Werner, J. Navaridas, M. Lujàn, Amon: An advanced mesh-like optical NoC, in: *23rd IEEE Annual Symposium on High-Performance Interconnects*, 2015, pp. 52–59.
- [31] N. Kirman, M. Kirman, R. K. Dokania, J. F. Martinez, A. B. Apsel, M. A. Watkins, D. H. Albonesi, Leveraging optical technology in future bus-based chip multiprocessors, in: *Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture*, 2006, pp. 492–503.
- [32] S. Le Beux, H. Li, G. Nicolescu, J. Trajkovic, I. O'Connor, Optical crossbars on chip, a comparative study based on worst-case losses, *Concurrency and Computation: Practice and Experience* 26 (15) (2014) 2492–2503.
- [33] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, A. Choudhary, Firefly: Illuminating future network-on-chip with nanophotonics, in: *Proceedings of the 36th Annual International Symposium on Computer Architecture, ISCA '09*, 2009, pp. 429–440.
- [34] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil, J. H. Ahn, Corona: System implications of emerging nanophotonic technology, in: *Proceedings of the 35th Annual International Symposium on Computer Architecture, ISCA '08*, 2008, pp. 153–164.
- [35] D. Vantrease, N. Binkert, R. Schreiber, M. H. Lipasti, Light speed arbitration and flow control for nanophotonic interconnects, in: *Microarchitecture, 42nd Annual IEEE/ACM International Symposium on (MICRO-42)*, 2009, pp. 304–315.
- [36] W. Fu, T. Chen, RCBUS: Row-column bus topology for optical network-on-chip, *Electronics and Electrical Engineering* 18 (8) (2012) 85–90.

- [37] J. Ouyang, Y. Xie, Enabling quality-of-service in nanophotonic network-on-chip, in: Proceedings of the 16th Asia and South Pacific Design Automation Conference, IEEE Press, 2011, pp. 351–356.
- [38] Y. Pan, J. Kim, G. Memik, Flexishare: Channel sharing for an energy-efficient nanophotonic crossbar, in: High Performance Computer Architecture (HPCA), IEEE 16th International Symposium on, IEEE, 2010, pp. 1–12.
- [39] C. Li, M. Browning, P. V. Gratz, S. Palermo, LumiNOC: A power-efficient, high-performance, photonic network-on-chip, Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on 33 (6) (2014) 826–838.
- [40] Y. Pan, J. Kim, G. Memik, FeatherWeight: low-cost optical arbitration with QoS support, in: Proceedings of the 44th Annual IEEE/ACM International Symposium on Microarchitecture, ACM, 2011, pp. 105–116.
- [41] A. Joshi, C. Batten, Y.-J. Kwon, S. Beamer, I. Shamim, K. Asanovic, V. Stojanovic, Silicon-photonic clos networks for global on-chip communication, in: Proceedings of the 3rd ACM/IEEE International Symposium on Networks-on-Chip, IEEE Computer Society, 2009, pp. 124–133.
- [42] J. Mo, J. Walrand, Fair end-to-end window-based congestion control, IEEE/ACM Transactions on Networking (ToN) 8 (5) (2000) 556–567.
- [43] D. P. Bertsekas, Nonlinear Programming, Athena Scientific, 1999.
- [44] S. Boyd, L. Vandenberghe, Convex Optimization, Cambridge University Press, 2003.
- [45] S. Koohi, S. Hessabi, All-optical wavelength-routed architecture for a power-efficient network on chip, Computers, IEEE Transactions on 63 (3) (2014) 777–792.
- [46] D. Zhao, Y. Wang, SD-MAC: Design and synthesis of a hardware-efficient collision-free QoS-aware MAC protocol for wireless network-on-chip, Computers, IEEE Transactions on 57 (9) (2008) 1230–1245.
- [47] B. Parhami, Computer Arithmetic: Algorithms and Hardware Designs, Oxford University Press, Oxford, UK, 2000.
- [48] M. Ercegovac, L. Imbert, D. Matula, J.-M. Muller, G. Wei, Improving Goldschmidt division, square root, and square root reciprocal, Computers, IEEE Transactions on 49 (7) (2000) 759–763.

- [49] Cacti 6.5, <http://www.hpl.hp.com/research/cacti>.
- [50] S. Li, J. H. Ahn, R. D. Strong, J. B. Brockman, D. M. Tullsen, N. P. Jouppi, The mcpat framework for multicore and manycore architectures: Simultaneously modeling power, area, and timing, *ACM Transactions on Architecture and Code Optimization (TACO)* 10 (1) (2013) 5.
- [51] F. Kashfi, S. Fakhraie, S. Safari, A 65nm 10GHz pipelined mac structure, in: *Circuits and Systems, 2008. ISCAS 2008. IEEE International Symposium on*, 2008, pp. 460–463.
- [52] J. Mukundan, J. F. Martinez, Morse: Multi-objective reconfigurable self-optimizing memory scheduler, in: *High Performance Computer Architecture (HPCA), 2012 IEEE 18th International Symposium on*, IEEE, 2012, pp. 1–12.
- [53] A. Varga, The OMNeT++ discrete event simulation system, in: *Proceedings of the European simulation multiconference*, 2001, pp. 319–324.
- [54] T. E. Carlson, W. Heirman, L. Eeckhout, Sniper: exploring the level of abstraction for scalable and accurate parallel multi-core simulation, in: *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, ACM, 2011, p. 52.